

National Computational Infrastructure for Lattice Gauge Theory

Principal Investigators

N. Christ (Columbia U.), M. Creutz (BNL), P. Mackenzie (Fermilab),
J. Negele (MIT), C. Rebbi (Boston U.), S. Sharpe (U. Washington),
R. Sugar (UCSB) and W. Watson, III (JLab)

December 18, 2001

Executive Summary

The long range objective of our collaboration is to create the hardware and software infrastructure needed for terascale simulations of quantum chromodynamics (QCD), the sector of the Standard Model of elementary particle physics that describes the strong interactions. Such simulations are necessary for understanding some of the most fundamental quantities in high energy and nuclear physics, thus supporting the Department of Energy's large experimental efforts in these fields. In our SciDAC proposal we requested funds for a three year software development and hardware prototyping effort. In the future we plan another proposal to the Department of Energy for a distributed topical computing center to provide the terascale platforms on which these simulations will be performed.

The twentieth century was an era of striking progress towards comprehending the fundamental structure of matter, beginning with the discovery of quantum mechanics and atomic physics, progressing to nuclear physics, and culminating with the Standard Model of elementary particle physics. However, traditional analytical tools have proven inadequate to extract many of the predictions of QCD. Our understanding of nature will remain fundamentally deficient until we know how the rich and complex structure of strongly interacting matter, which comprises most of the known mass of the universe, arises from the interactions among quarks and gluons.

At present, the only method to extract non-perturbative predictions of QCD from first principles and with controlled systematic errors is through large scale numerical simulations of lattice gauge theory. Recent refinements of numerical algorithms coupled with major increases in the capabilities of massively parallel computers have brought these simulations to a new level. It is now possible to calculate a few crucial quantities to an accuracy comparable with their experimental determination. The strong coupling constant and the masses of the c and b quarks are notable examples. Furthermore, the experience we have gained allows confident predictions for the computing resources required for accurate determinations of a broad range of fundamental quantities. Many of the most important calculations which we propose to carry out will require long term use of multi-teraflops computing facilities.

United States physicists invented lattice gauge theory and play a leadership role in the field. To maintain our status and provide needed support for the U.S. high energy and nuclear experimental programs, we must move quickly to construct the substantial computational infrastructure required for terascale simulations. Unfortunately, the U.S. is falling behind Europe and Japan in building such infrastructure. Our SciDAC effort represents a first step towards a coherent national plan in support of the next level of scientific discovery in lattice gauge theory. The envisioned terascale computing resources will enable calculations essential to precision tests of the Standard Model, to understand the structure of nucleons and other hadrons, and to determine the properties of hadronic matter under extreme conditions.

Massively parallel computers are ideally suited for lattice gauge calculations, having been successfully exploited for many years. Furthermore, it has proven far more cost effective for lattice gauge theorists to build their own computers than to make use of general purpose supercomputers. In doing so they can take optimal advantage of simplifying features of lattice QCD, such as regular grids and uniform, predictable communications. General purpose machines must perform well for a wide variety of problems, including those requiring irregular or adaptive grids, non-uniform communication patterns, and massive input/output capabilities. Thus, commercial supercomputers require considerably more expensive communication systems than are needed for lattice QCD. For this reason, our plans include the development of both hardware and software infrastructure, pursuing both customized clusters of commodity components and the development of a fully custom lattice computer. The concept of a topical computing facility as set out in the Office of Science's computing plan is particularly well suited for lattice gauge theory.

The objective of our SciDAC project is to support the software development necessary to productively

exploit the envisioned multi-teraflops computing facilities. Two architectures will be targeted: large commodity clusters and the QCDOC, the next generation of the highly successful Columbia/ Riken/ BNL special purpose computers.

A flexible, user-friendly software environment is critical to the development of efficient new algorithms and computational methods. We will collaborate closely with our colleagues in computer science, computer engineering and applied mathematics on the software and algorithms research and the hardware prototyping needed for terascale QCD systems. The interplay between hardware, software, and algorithms research is central to our work.

The software infrastructure we are developing will enable members of the national lattice gauge community to achieve high performance on future terascale systems, while focusing their efforts on frontier questions in physics. This will require a variety of components for high application performance, including standardized libraries with highly tuned code for common computationally intensive tasks and optimized communications primitives, implemented for each target architecture. At a higher level, the project will deliver standard programming models, portable applications, and user-friendly interfaces. The software infrastructure will also incorporate components necessary to schedule and monitor jobs, standards for data formats, and other tools to support wide access and management of large computational and data resources.

Scalability of lattice applications to large clusters of commodity computers, including symmetric multiprocessors (SMPs), will be a key development area. Systems containing hundreds of processors are required to provide adequate test-beds for the software and for the viability of multi-teraflops clusters for QCD. To this end we are expanding the clusters currently under development at Fermilab and by the MIT/Jefferson Lab consortium. These systems will be made available to lattice researchers for their scientific work and to further test and “harden” the software infrastructure. In this activity, the clusters will be available to the full U.S. lattice gauge theory community, as will all of the hardware and software infrastructure created under this project.

This work will provide the basis for a future proposal requesting distributed topical computing resources for the study of QCD. To meet the full scientific challenges of QCD, and to allow U.S. lattice gauge theorists to compete effectively with researchers in Europe and Japan, we plan to propose the coordinated development of three open national facilities of 10 Tflops scale in the period 2003-2007. The first terascale facility will be a QCDOC based machine to be located at Brookhaven National Laboratory, at a cost of less than \$1 per sustained Mflops. The architectures used for the 10 Tflops facilities at Fermilab and Jefferson Lab will be chosen on the basis of the experience gained with clusters and the QCDOC, on their cost effectiveness, and on their ability to support the scientific program we have set out. This plan will optimally position us to exploit future technologies for fundamental physics calculations, and this multi-pronged approach will be crucial to maintaining flexibility in the future.

Information regarding current activities and progress can be found on the project’s webpage:
<http://physics.bu.edu/~brower/>.

We intend to deploy grid tools partially developed by the Particle Physics Data Grid Collaboratory to help us make the hardware test beds at Jefferson Lab and Fermilab more accessible to our community, and to test their usefulness for a future distributed Lattice Computing Facility.